

[Previous](#) [Next](#) [Contents](#) [Index](#)

Chapter 13 Managing Replication

Replication is an important mechanism for extending your directory service beyond a single server configuration. In the following sections, this chapter describes how you can use replication for your directory service:

- ["Replication Overview"](#)
- ["Managing Supplier-Initiated Replication \(SIR\)"](#)
- ["Managing Consumer-Initiated Replication \(CIR\)"](#)
- ["Initializing Consumers"](#)
- ["Monitoring Replication Status"](#)
- ["Replication Algorithms"](#)
- ["Machine data"](#)

For conceptual information on how you can use replication in your directory service, see the Netscape Directory Server Deployment Guide.

● Replication Overview

Replication is the mechanism by which directory data is automatically copied from one directory server to another. Using replication, you can copy entire directory trees or subtrees between servers. Updates of any kind—entry additions, modifications, or even deletions—are automatically mirrored to other directory servers using replication.

The master copy of all directory data is stored in one and only one directory location. The server that controls this master copy of directory data is called the supplier server. Only the supplier server can modify or delete data in the master copy. The supplier server can then propagate changes made to its directory data to other servers known as consumer servers.

There are two basic forms of replication available to you: supplier-initiated replication and consumer-initiated replication. Supplier-initiated replication allows you to configure a supplier server to push data to one or more consumer servers. Consumer-initiated replication allows you to configure consumer servers to pull directory data from a supplier server.

Directory servers use replication agreements to define replication. A replication agreement identifies the directory objects to replicate, the times during which replication can occur, and the server to which the replicated data is pushed or the server from which replicated data is pulled.

This chapter provides detailed information on how to set up replication agreements. For more overview information on what replication is and how you might use it, see the *Netscape Directory Server Deployment Manual*.

● Managing Supplier-Initiated Replication (SIR)

This section provides information on how you can manage SIR agreements in the following sections:

- "Configuring Servers for SIR"
- "Creating an SIR Agreement"
- "Duplicating an SIR Agreement"
- "Editing an SIR Agreement"

For more information about how to use replication within your enterprise, see the Netscape Directory Server Deployment Guide.

Configuring Servers for SIR

Before you can create an SIR agreement, you have to configure your servers to be either a supplier or a consumer. To do this, you must configure basic information about the servers.

To configure servers for supplier-initiated replication you need to do the following:

- On the consumer server—Configure a DN and a corresponding password for the supplier to use to bind to the consumer.

If you want communication between the consumer and supplier to take place over SSL you also need to configure a certificate subject DN on the consumer server. For specific instructions, see "Configuring the Supplier DN for SIR".

- On the supplier server—Configure a change log directory and then restart the supplier server. For specific instructions, see "Configuring the Change Log for SIR".

The change log is a special directory maintained by the supplier server that identifies the changes made to the server's primary directory tree.

Once you have set up the consumer and supplier servers, you are ready to create

the agreement. For specific instructions, see "[Creating an SIR Agreement](#)".

Configuring the Supplier DN for SIR

Before you configure an SIR agreement, you need to configure a supplier DN and password on your consumer server. The supplier server uses this DN to bind to the consumer server during replication. The Supplier DN is a special distinguished name that does not actually exist in your directory tree. Instead, it is identified by the [Supplier DN](#) parameter in the `slapd.conf` file.

Entries supplied to the consumer server from another server can only be updated if the LDAP client binds as the supplier DN; all other update operations for the supplied data will be referred to the supplier server that masters the directory data. It is therefore important that the only LDAP clients that bind to this server using the supplier DN are supplier servers.

You can configure the consumer server to accept simple or password-based bind operations from the supplier, and you can configure the server to accept certificate-based authentication. Certificate-based authentication is more secure because the password used to authenticate is encrypted and does not need to be stored in cleartext on the supplier server. To configure the supplier DN and authentication method:

1. On the consumer server's Directory Server Console, select the Configuration tab and then select the Replication Agreements folder.
2. Select the Consumer Server Settings tab in the right pane.
3. In order to use simple authentication or certificate-based authentication, you must enter the DN and password you want the supplier server to use to bind to this server in the Supplier DN, and Supplier password text boxes.
4. If you want to use certificate-based authentication, type or paste the subject DN of the certificate that the supplier server will use to bind to this server in the Supplier Certificate Subject DN box.

If you have more than one server supplying entries to this consumer, enter a subject DN for each supplier server. Each DN should be placed on a separate line in the box.

You can find the subject DN of the certificate used by a supplier server from the supplier server's console. For specific instructions on how to do this, see "Setting up Encryption Security" in *Managing Servers with Netscape Console*.

Configuring the Change Log for SIR

Before a server can supply directory entries to consumer servers, you must configure a change log on the supplier server. The change log is a special database maintained by the supplier server that identifies the changes made to the server's primary directory tree. To configure the change log:

1. On the supplier server's Directory Server Console, select the Configuration

tab and then select the Replication Agreements folder.

2. Select the Supplier Server Settings tab in the right pane.
3. Type the full path to the directory where you want the server to store the change log in the Changelog Database Directory text box.

This directory must be located on the supplier's local disk. If you want the directory server to suggest a pathname, click Use Default.

4. Enter a DN to use as the change log's directory suffix in the Changelog Suffix text box. Typically, this suffix is: cn=changelog.
5. Either enter the maximum number of records you want the change log to record in the Max Changelog Records text box, or if you do not want to set a maximum number of entries for the change log, select Unlimited.
6. If you want the server to remove entries from the change log after they reach a certain age, specify that age in seconds, minutes, hours, days, or weeks in the Max Changelog Age fields.

If you do not want to configure a maximum age, select Unlimited; the server will not remove entries from the change log due to their age.

7. Click Save.
8. Restart the directory server.

You are now ready to configure an SIR agreement.

Creating an SIR Agreement

To create a supplier-initiated replication agreement:

1. On the Directory Server Console, select the Configuration tab.
2. Right-click the Replication Agreements folder and select New Replication Agreement.

The Replication Agreement Wizard appears. This wizard takes you through the steps of setting up a replication agreement.

3. On the wizard dialog, select Supplier Initiated Agreement and click Next.

A dialog box appears that allows you to provide a name for the replication agreement.

4. Provide a name for the replication agreement and click Next.

The Replication Agreement form displays.

5. Select a consumer from the Consumer drop-down menu or, click Other to manually enter the host and port number of the consumer server.
6. If you want the servers to use SSL during replication, select the "Using encrypted SSL connection" checkbox.
7. If you want the servers to use SSL client authentication, select SSL Client Authentication.

You cannot use SSL Client Authentication unless you have specified that the server use encrypted SSL connections in Step 6.

8. If you want the servers to communicate using simple authentication (with or without SSL), select Simple Authentication and then provide the Bind DN and password.
9. Enter the subtree you want replicated in the Subtree text box or, choose Browse to select the node you want to replicate.
10. When you are finished, click Next. This brings up the Replication Schedule form.
11. If you do not want to limit replication to explicit time periods, select "Always keep directories in sync". Alternatively, you can identify the time of day and day(s) of the week when replication can occur by selecting "Sync on the following days".

If you choose to limit replication to specific days and times, select the checkbox next to the day(s) and enter the hours between which replication can take place. Any replication activity occurring when the specified time interval ends will be completed, but no new replication processes will be started outside the specified replication interval.

When you are finished scheduling the replication agreement, click Next. The Consumer Initialization dialog displays.

12. Choose one of the following options:
 - o Do Not Initialize Consumer—Select this option if you do not want the consumer initialized automatically or the LDIF file created.
 - o Initialize Consumer Now—Select this if you want the server to initialize the consumer when you finish creating the replication agreement. For performance reasons, this is not recommended for databases larger than 10,000 entries.
 - o Create Consumer Initialization File—Select this if you want the server to export the replicated tree to LDIF so you can manually import it to the consumer. If you choose to have the server export to LDIF, supply the LDIF filename in the field provided.
13. Click Next.

The summary dialog appears.

You need to initialize the consumer before replication can occur. If you choose not to initialize the consumer now, see ["Initializing Consumers"](#) for instructions on how to do it later.

14. Make sure that the information on the summary dialog box is correct.

If any information is incorrect, click Back to step back through the dialogs and change the information. When you are finished, click Done. The server creates the replication agreement and dismisses the replication wizard.

Duplicating an SIR Agreement

To add a consumer to a supplier-initiated agreement you need to duplicate the SIR Agreement and then update the consumer setting to point to the new consumer. In reality, you are not adding a consumer to an existing agreement, so much as copying the details of an existing agreement and adding the new consumer to the copy. To do this:

1. On the Directory Server Console, select the Configuration tab.
2. Open the Replication Agreements folder and then the Supplier Initiated Agreements folder.
3. Right-click the replication agreement in the tree and select Duplicate Replication Agreement from the pop-up menu.

The Replication Agreement Wizard appears with the original settings displayed in the dialog boxes.

4. Enter a unique name representative of this agreement and click Next.
5. On the dialog that appears, enter the new consumer in the Consumer text box.
6. Complete the forms and click Done when you are finished.

Editing an SIR Agreement

You can make changes to existing replication agreements using the Directory Server Console. To do this:

1. On the supplier server's Directory Server Console, select the Configuration tab.
2. Open the Replication Agreements folder and then the Supplier Initiated Agreements folder.
3. Select the replication agreement you want to edit.

4. To change the name of the replication agreement, select the Summary tab in the right pane and enter your changes in the Agreement Name text box.
5. To edit the scheduling information for the replication agreement, select the Schedule tab in the right pane.

If you do not want to limit replication to explicit time periods, select "Always keep directories in sync". Alternatively, you can identify the time of day and day(s) of the week when replication can occur by selecting "Sync on the following days". If you choose to limit replication to specific days and times, select the checkbox next to the day(s) and enter the hours between which replication can take place. Any replication activity occurring when the specified time interval ends will be completed, but no new replication processes will be started outside the specified replication interval.

6. To change the general settings for this replication agreement, such as the consumer server you want to replicate to, what you want replicated, and whether or not SSL is used for the connection, select the Content tab in the right pane.
 - o To change the consumer this supplier replicates to, select a different consumer from the Consumer drop-down menu or, click Other to manually enter the host and port number of the new consumer server.
 - o If you want the servers to use SSL during replication, select the "Using encrypted SSL connection" checkbox.
 - o If you want the servers to use SSL client authentication, select SSL Client Authentication. You cannot use SSL Client Authentication unless you have specified that the server use encrypted SSL connections.

In order for you to select this option, you must first configure SSL for both your supplier and consumer servers (see [Chapter 11, "Managing SSL"](#)), and configure your consumer server to recognize the subject DN your supplier server's certificate as the supplier DN (see ["Configuring the Supplier DN for SIR"](#)).

- o If you want the servers to communicate using simple authentication (with or without SSL), select Simple Authentication and then provide the Bind DN and password.
- o You can change the subtree you want replicated in the Content Replicate text box. Alternatively, choose Browse to browse the contents of the supplier server.

If you are going to replicate a subtree, you must make sure the appropriate parent entry is available on the consumer server. For example, if you are replicating the ou=people, o=airius.com subtree, then you must first make sure the consumer server contains the o=airius.com entry.

7. When you have finished making changes, click Save.

● Managing Consumer-Initiated Replication (CIR)

This section provides information on managing CIR agreements in the following sections:

- "Configuring Servers for CIR"
- "Creating a CIR Agreement"
- "Editing a CIR Agreement"

For more information about how to use replication within your enterprise, see the Netscape Directory Server Deployment Guide.

Configuring Servers for CIR

Before you can create a CIR agreement, you have to configure your servers to be either a supplier or a consumer. To do this, you must configure basic information about the servers.

To configure servers for consumer-initiated replication you need to:

- Configure a change log directory on the supplier server and then restart the supplier server. For specific instructions, see "Configuring the Change Log for CIR".
- Set up consumer access to the change log directory on the supplier server. This includes creating a DN for the consumer to use to connect to the supplier server and setting the access control instructions for the DN to allow the consumer to search and read the change log. For specific instructions, see "Providing Consumer Access to the Change Log for CIR".

Once you have set up the consumer and supplier servers, you are ready to create the agreement as described in "Creating a CIR Agreement".

Configuring the Change Log for CIR

Before the consumer server can collect updated information from the supplier server, you must configure a change log for the supplier server. The change log is a special database maintained by the supplier server that identifies the changes made to the server's primary directory tree.

To configure the change log:

1. On the supplier server's Directory Server Console, select the Configuration tab and then select the Replication Agreements folder.
2. Select the Supplier Server Settings tab in the right pane.

3. In the Changelog Database Directory text box, type the full path to the directory where you want the server to store the change log.

This directory must be located on the supplier's local disk. If you want the directory server to suggest a pathname, click Use Default.

4. In the Changelog Suffix text box, enter a DN to be used as the change log's directory suffix. Typically, this suffix is: cn=changelog.
5. Either enter the maximum number of records you want the change log to record in the Max Changelog Records text box, or if you do not want to set a maximum number of entries for the change log, select Unlimited.
6. If you want the server to remove entries from the change log after they reach a certain age, specify that age in seconds, minutes, hours, days, or weeks in the Max Changelog Age fields. If you do not want to configure a maximum age, select Unlimited; the server will not remove entries from the change log due to their age.
7. Click Save.
8. Restart the directory server.

Providing Consumer Access to the Change Log for CIR

Before you can use a consumer-initiated replication agreement, the consumer server must be able to read the change log directory.

To provide consumer access to the change log for CIR:

1. Create a directory entry (pseudo-user) that can be used by your consumer servers to read your change log.

This directory entry does not have to be created in your change log directory tree; it can be a normal entry in your primary directory tree provided the entry contains the userPassword attribute. For information, see Chapter 9, "Managing Directory Entries."

2. At the root level of your change log tree, create an ACI statement that grants the user from step 1 read, search, and compare access to the entire change log tree. In addition, this ACI should grant full read, search, and compare privileges for the tree or subtree that the consumer server will retrieve from the supplier server. For more information, see "Configuring the Change Log for CIR" and Chapter 5, "Managing Access Control."

For security reasons, Netscape recommends that you do not configure anonymous access for your change log directory tree. Also, you should grant only read, search, and compare access to the DN with which your consumers bind to your supplier; do not provide any form of write or delete access to this tree.

Finally, for logging and auditing purposes, you may want to configure a different directory entry (pseudo-user) for each consumer server. This allows you to track which consumer is binding to your server and when. For information on tracking

access, see ["Viewing the Access Log"](#).

Creating a CIR Agreement

To create a consumer-initiated replication agreement:

1. On the consumer server's Directory Server Console, select the Configuration tab.
2. Right-click the Replication Agreements folder and select New Replication Agreement.

This brings up the Replication Agreement Wizard which takes you through the steps of setting up a replication agreement.

3. On the wizard dialog box, select Consumer Initiated Agreement and click Next.

This displays a dialog that allows you to provide a name for the replication agreement.

4. Enter a name for the replication agreement and click Next.

This displays the Replication Agreement dialog box.

5. Select the supplier server from which you want the consumer to pull replicated information or, click Other to manually enter the host and port number of the supplier server.
6. If you want the servers to use SSL during replication, select the "Using encrypted SSL connection" checkbox.
7. If you want the servers to use SSL client authentication, select SSL Client Authentication. You cannot use SSL Client Authentication unless you have specified that the server use encrypted SSL connections in [Step 6](#).
8. If you want the servers to communicate using simple authentication (with or without SSL), select Simple Authentication and then provide the Bind DN and password.
9. Enter the subtree you want to replicate in the Content Replicate text box. Make sure the parent of the subtree exists on the consumer server. You can also click Browse to browse the contents of the supplier server. When you are finished, click Next. The Replication Schedule dialog box appears.
10. You must configure the consumer server to periodically check the supplier server to see if there are any pending updates by entering a time interval in the Update Interval text box. The interval is defined in minutes.
11. If you do not want to limit replication to explicit time periods, select "Always keep directories in sync". Alternatively, you can identify the time of day and day of week when replication can occur by selecting "Sync on the following

days". If you choose to limit replication to specific days and times, select the checkbox next to the day(s) and enter the hours between which replication can take place. Any replication activity occurring when the specified time interval ends will be completed, but no new replication processes will be started outside the specified replication interval. When you are finished scheduling the replication agreement, click Next.

12. Choose one of the following options:

- o Do Not Initialize Consumer—Select this option if you do not want the consumer initialized automatically or an LDIF file created.
- o Initialize Consumer Now—Select this if you want the server to initialize the consumer when you finish creating the replication agreement. This is not recommended for databases larger than 10,000 entries.

13. Click Next. The summary dialog appears.

You need to initialize the consumer before replication can occur. If you choose not to initialize the consumer now, see "Initializing Consumers" for instructions on how to do it later.

14. Make sure that the information in the dialog box is correct. If any information is incorrect, click Back to step back through the dialogs and change the information. When you are finished, click Done. The server creates the replication agreement and dismisses the replication wizard.

Duplicating a CIR Agreement

To add a supplier to a consumer-initiated agreement you need to duplicate the CIR Agreement and then update the supplier setting to point to the new supplier. In reality, you are not adding a supplier to an existing agreement, so much as copying the details of an existing agreement and adding the new supplier to the copy. To do this:

1. On the Directory Server Console, select the Configuration tab.
2. Open the Replication Agreements folder and then the Consumer Initiated Agreements folder.
3. Right-click the replication agreement in the tree and select Duplicate Replication Agreement from the pop-up menu. This brings up the Replication Agreement Wizard with the original settings displayed in the dialog boxes.
4. Enter a unique name representative of this agreement and click Next.
5. On the dialog that appears, enter the new supplier in the Supplier text box.
6. Complete the rest of forms and click Done when you are finished.

Editing a CIR Agreement

You can make changes to existing replication agreements using the Directory Server Console. To do this:

1. On the consumer server's Directory Server Console, select the Configuration tab.
2. Open the Replication Agreements folder and then the Consumer Initiated Agreements folder.
3. Select the replication agreement you want to edit.
4. To change the name of the replication agreement, select the Summary tab in the right pane and enter your changes in the Agreement Name text box.
5. To edit the scheduling information for the replication agreement, select the Schedule tab in the right pane.
 1. You must configure the consumer server to check the supplier server to see if there are any pending updates. To do this, enter a time interval in the Update Interval text box.
 2. If you do not want to limit replication to explicit time periods, select "Always keep directories in sync". Alternatively, you can identify the time of day and day of week when replication can occur by selecting "Sync on the following days". If you choose to limit replication to specific days and times, select the checkbox next to the day(s) and enter the hours between which replication can take place. Any replication activity occurring when the specified time interval ends will be completed, but no new replication processes will be started outside the specified replication interval.
6. To change the general settings for this replication agreement, such as the supplier server you want to replicate from, what you want replicated, and whether or not SSL is used for the connection, select the Content tab in the right pane.
 1. To change the supplier from which this consumer gets information, select a different supplier from the Supplier drop-down menu. Alternatively, click Other to manually enter the host and port number of the new supplier server.
 2. If you want the servers to use SSL during replication, select the "Using encrypted SSL connection" checkbox.
 3. If you want the servers to use SSL client authentication, select SSL Client Authentication. You cannot use SSL Client Authentication unless you have specified that the server use encrypted SSL connections.

In order for you to select this option, you must first configure SSL for both your supplier and consumer servers (see [Chapter 11, "Managing SSL"](#)), and configure your consumer server to recognize the subject DN or your supplier server's certificate as the supplier DN (see ["Configuring the Supplier DN for SIR"](#)).

7. If you want the servers to communicate using simple authentication (with or without SSL), select Simple Authentication and then provide the Bind DN and password.
8. You can change the subtree you want replicated in the Content Replicate text box. Alternatively, choose Browse to browse the contents of the supplier server.
9. When you have finished making changes, click Save.

● Removing the Change Log

If you change the change log suffix, the entries in the change log are no longer valid. For this reason, you need to remove the change log and create a new one. If you remove the change log, you will need to reinitialize your consumer servers. You can remove the change log using the Directory Server Console. To do this:

1. On the supplier server's Directory Server Console, select the Configuration tab.
2. Select the Replication Agreements folder in the navigation tree in the left pane and then the Supplier Server Settings tab in the right pane.
3. Click Remove Changelog.
4. Click Save.
5. Restart the directory server.
6. Reinitialize your consumer servers.

● Initializing Consumers

There are two ways that you can initialize a consumer:

- Online Consumer Creation—This method is the easiest to perform but is prohibitively time consuming for databases that are larger than 5,000 - 10,000 entries in size.
- Manual Consumer Creation—This is the more difficult but most efficient method.

This section first describes consumer initialization in detail and then provides instructions on both consumer creation methods.

When to Initialize a Consumer

After you have created a replication agreement, you must initialize the consumer. That is, you must physically copy directory data from the supplier server to the

consumer server so that future changes can be replayed to consumer servers.

Consumer initialization involves copying the replicated directory entries from the supplier server to the consumer server. When these entries are placed on the consumer server, the appropriate `copiedFrom` attribute must also be placed on the replicated tree (see "Replication Algorithms" for details).

Once the tree has been physically placed on the consumer, the supplier server can begin replaying update operations to the consumer server (SIR) or the consumer can begin requesting data from the supplier (CIR).

In addition, any attempts to modify data on the consumer that is owned by the supplier are referred to the supplier server. For more information about referrals, see Chapter 14, "Managing Referrals."

Under normal operations, the consumer should not ever have to be initialized again. However, there are several major events that can require a reinitialization of the consumer server:

- The supplier's database version number does not match the version number stored on the consumer for the replicated entries. This will happen if the supplier's database is either reloaded from a backup.

The database version number is a unique identifier that allows the supplier and the consumer to know that the database has not been reloaded since the last synchronization.

- The change log on the supplier server is damaged to the extent that the supplier cannot determine what changes to replay to the consumer. This can happen if the supplier's change log becomes corrupted (such as might happen in the event of a disk failure) or if the change log is trimmed before the trimmed changes can be replayed to the consumer server.

This situation is most likely to arise if the consumer's database is restored from a backup and the supplier's change log was truncated sometime after that backup was taken.

Change logs are trimmed based on the Maximum Changelog Age and Maximum Changelog Size parameters. For more information, see "Configuring the Change Log for CIR."

- The supplier server is continually repairing inconsistencies on the consumer server, or the supplier is unable to repair data inconsistencies on the consumer server.

In almost all cases, once you have initialized the consumer server, the supplier can successfully repair inconsistencies on a consumer server. Further, even if it cannot repair data inconsistencies, the supplier will continue to replicate to the consumer server. When the supplier server detects a data inconsistency on a consumer server, the supplier issues the following message to the error log:

Inconsistency detected while replaying change <n>,

entry <DN>, to replica <host>:<port>/<DN>

This message also indicates whether the inconsistency could be repaired.

The process that you use to initialize or reinitialize a consumer differs depending on the type of consumer creation you use. See "[Online Consumer Creation](#)" (next) or "[Manual Consumer Creation](#)" for more information.

Note. When a consumer server is being initialized via online consumer initialization, all operations (including searches) on the supplied tree are referred to the supplier server until the initialization process is completed.

Online Consumer Creation

Online consumer creation is the easiest way to initialize or reinitialize a consumer. However, this process can be very time consuming, and for large databases you may find that manual consumer creation is a more appropriate approach (refer to "[Manual Consumer Creation](#)" for more information).

Online consumer creation works by moving data from the supplier to the consumer server over LDAP. That is, the replicated information is placed on the consumer server using LDAP add operations.

Before using online consumer creation, consider the performance implications of this method of consumer initialization. On a reasonably fast single processor (such as an Intel Pentium II or a Sun Sparc Ultra 1), you can expect online consumer creation to proceed at the following rates:

- For fresh initializations (that is, if the consumer server's tree is empty), the supplier can add 9,000 to 36,000 entries per hour. The actual rate will depend on characteristics of your server such as the size of your entries, the amount of indexing your consumer is performing, the speed of your disk, the amount of RAM available to the consumer server, and the speed of your networks.
- For reinitializations, the supplier can add from 4,500 to 18,000 entries per hour. The performance drops by half because the online consumer creation process deletes all previously replicated entries from the consumer server before the consumer is initialized with a fresh set of data.

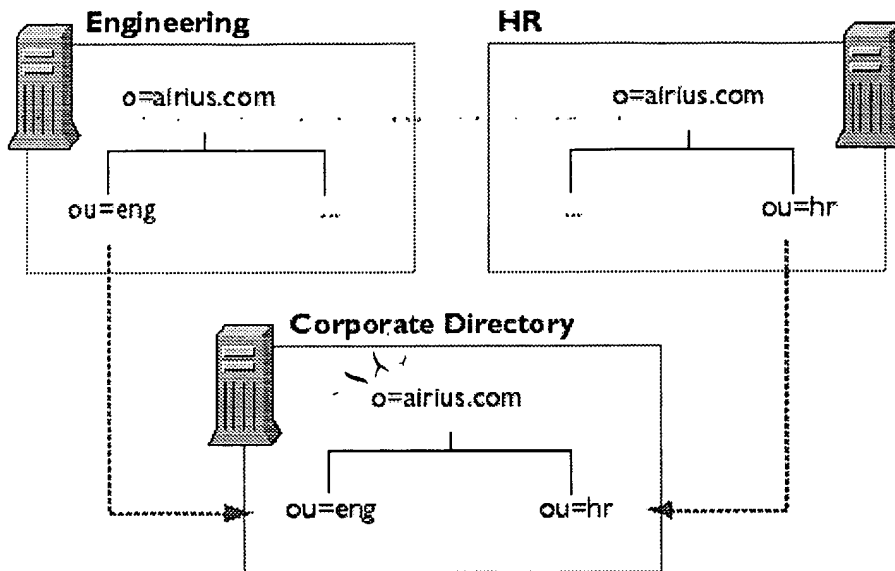
When You Should Use Online Consumer Creation

Essentially, you should always use online consumer creation unless you find the time that it takes to complete this operation objectionable.

You should always use online consumer creation when the consumer server has multiple subtrees supplied by different supplier servers, as shown in the figure below. In this case, you cannot use manual consumer initialization to initialize the consumer because the import process used in manual consumer initialization replaces the entire database, including the subtrees supplied by other servers.

In the following figure, both Engineering and HR are supplier servers that replicate

data to the main corporate directory.



How to Use Online Consumer Creation

To use online consumer creation:

1. Create a replication agreement. For details on creating replication agreements, see ["Creating an SIR Agreement"](#) or ["Creating a CIR Agreement"](#).
2. On the supplier server's Directory Server Console, right-click the appropriate replication agreement on the Configuration tab and select **Initialize Consumer** from the pop-up menu.
3. Click **Yes** in the confirmation box.

Online consumer creation begins immediately. You can check the status of the online consumer creation on the console's Status tab. For more information about monitoring replication and initialization status, see ["Monitoring Replication Status"](#). If online consumer creation is in progress, the status shows that a replica is being initialized. To update this window, click **Refresh**. When online consumer creation finishes, the status changes to reflect this.

You can configure your server to automatically reinitialize a consumer when the server detects an unexplainable inconsistency. To do this, place the following line in the `slapd.conf` file of either the supplier server (for supplier-initiated replication) or the consumer server (for consumer-initiated replication):

```
orcauto on
```

See ["Directory Server Configuration Files"](#) for information on where the server stores `slapd.conf`.

The ["Enable Online Consumer Creation"](#) `slapd.conf` parameter causes the

consumer to be reinitialized if a version number mismatch occurs between the supplier server's database and the replicated entries, or if the supplier is unable to replay changes to the consumer due to problems with the change log (see ["When to Initialize a Consumer"](#) for more information).

Manual Consumer Creation

Manual consumer creation is the fastest method of consumer initialization for sites that are replicating very large numbers of entries. However, the manual process is more complicated than the online creation process.

You should use the manual process whenever you find that the online process is inappropriate due to performance concerns. However, you should never use the manual process if your consumer server contains directory data that is mastered by more than one directory server. That is, use the manual process only if your entire consumer server's database is supplied by a single supplier server. This is because typically when you use `ldif2db`, the database is completely overwritten. (The exception to this is the configuration tree `o=NetscapeRoot` which you can choose not to overwrite during initialization.)

For information on consumer initialization, see ["Initializing Consumers"](#). For information on the online consumer creation process, see ["Online Consumer Creation"](#).

To manually initialize or reinitialize a server:

1. Create a replication agreement as described in ["Creating an SIR Agreement"](#) or ["Creating a CIR Agreement"](#). When prompted, select the "Create consumer initialization file" radio button.
2. Import the LDIF file to the consumer server. See ["Importing the LDIF File to the Consumer Server"](#) for instructions.

Converting the Supplier Tree to LDIF

You can convert the tree to LDIF when you create a replication agreement by selecting "Create consumer initialization file" on the Consumer Initialization dialog in the Replication wizard. (See ["Creating an SIR Agreement"](#) or ["Creating a CIR Agreement"](#) for information.)

If you choose not to export the tree at that time, you can:

- right-click the replication agreement in the Directory Server Console and select "Create LDIF File" from the pop-up menu, or
- use the export command as described in ["Exporting Databases to LDIF"](#). If you convert the database to LDIF using the command-line, you must specify the `-r` argument to ensure that the `copiedFrom` attribute is included in the output.

Importing the LDIF File to the Consumer Server

Create your consumer server's database from the LDIF file by using either the Import command from the server console, or the `ns-slapd ldif2ldb` command-line utility. For more information, see ["Importing LDIF From the Server Console"](#), or ["Importing LDIF From the Command Line"](#).

If your consumer server contains data that is also mastered either by itself or by some other supplier server, then use the online consumer creation process (for details, see ["Online Consumer Creation"](#)). While it is possible to manually import this LDIF file, you must do so using `ldapmodify` which offers no performance improvement over online consumer creation because both mechanisms add entries over LDAP.

If you decide that you must manually initialize a consumer that contains data mastered by some other server than your supplier server, make sure you do the following:

- Create any entries that are parents of the replicated subtree on the consumer server before adding the replicated data. That is, if you are replicating `l=Minneapolis, ou=Global, o=airius.com`, make sure that you have created `ou=Global, o=airius.com` and `o=airius.com` on your consumer server before running the add operation.
- If you are reinitializing a consumer server, make sure you delete all of the contents of the replicated tree before running the add operation. If you do not delete the currently existing replicated tree, the server will fail the add operations, stating that the entries already exist.
- When you are adding or deleting replicated entries, bind to your consumer server using the supplier DN configured for that server. If you use any other DN (including the root DN), the consumer server will simply refer the modify operation to the supplier server.

Monitoring Replication Status

You can monitor replication status using the Directory Server Console.

To view a summary of replication status:

1. On the Directory Server Console, select the Status tab and then select Replication Agreements in the navigation tree in the left pane.

In the right pane, a table appears that contains information about each of the replication agreements configured for this server.
2. Click Refresh to update the contents of the tab.

The status information displayed is described in [Table 13.1](#).

Table 13.1 Directory Server Console - Status tab

Table	Description
-------	-------------

Header	
Agreement	Contains the name you provided when you set up the replication agreement. A red bullet to the left of the name indicates an error has occurred and replication cannot take place. A green bullet indicates that replication is occurring normally. A yellow bullet indicates that all of the changes have not yet been sent to the consumer; this does not always indicate an error condition.
Supplier	Specifies the supplier server in the agreement.
Consumer	Specifies the consumer server in the agreement.
Change-Number	<p>Indicates the last change number replayed to the consumer and the last change number available in the supplier's change log. For example: [7] - [10]</p> <p>"Unknown" indicates that the server has encountered an error and replication cannot continue or the server could not read one of the following:</p> <ul style="list-style-type: none"> • The last change number from the supplier • The copiedFrom on the consumer <p>These situations are normal if no changes have occurred on the supplier or if the consumer has not been initialized.</p>
Status	<p>Specifies the current state of the agreement. The possible values include:</p> <ul style="list-style-type: none"> • Idle—No replication is currently taking place through this agreement. This might appear if there are no changes to replay or if the replication agreement is not scheduled to start until a later time. • Synchronizing—Changes are currently being sent to the consumer. • Populating—Online Replica Creation is in progress; the consumer is being initialized. • Halted—the synchronization process has encountered an error and quit.

Replication Algorithms

This section describes the replication processes in detail for both supplier-initiated and consumer-initiated replication.

SIR Algorithm

If you are using supplier-initiated replication, it is the responsibility of the supplier server to determine when its consumer servers need to be updated. This process occurs as follows:

1. Based on a schedule that you set, the supplier server determines that it is time to synchronize a consumer. The directory server identifies those subtrees that are replicated and the servers to which it is supplying those trees by using directory entries contained in the Machine data tree. Each consumer server is identified by a separate replication agreement, stored beneath the machine data entry. Part of each replication agreement is an identification of the root point of the replicated tree and a schedule indicating when the consumer should be updated.
2. If there are changes to replay, the supplier server binds to the consumer server using the supplier DN and password that you provide. The supplier must use this special DN for the bind or all updates to the replicated tree will be referred back to the supplier server.

You use the Replication Agreements tabs to configure the supplier DN for a consumer server. For information, see ["Configuring the Supplier DN for SIR"](#).

3. Each replicated tree contained on a consumer server includes a `copiedFrom` attribute that identifies the supplier of the subtree. This attribute is maintained by the replication subsystem. The supplier server examines the `copiedFrom` attribute in the replicated subtree's root point to ensure that no other supplier is identified by this attribute and to determine if replication can occur.

If no such attribute exists for the subtree, the supplier server

4. Writes the attribute to the replicated root point along with the appropriate attribute value
5. Aborts and immediately retries the synchronization

The syntax for the `copiedFrom` attribute is as follows:

```
copiedFrom: host:port generationID last_change
```

For example:

```
copiedFrom: dir.airius.com:389 019980610154028 12
```

where `host:port` is the host name and port number of the supplier server, `generationID` is a timestamp generated for the supplier database when the database is created, and `last_change` is a number generated by the supplier server that increases by one for every change replayed to the consumer. The `generationID` and `last_change` number are only reset if the database is reloaded. Do not modify the `generationID` or the `last_change` number manually. If you do, replication will fail and you will have to reinitialize your consumer servers.

Before replication can occur, the supplier checks the consumer's `copiedFrom` attribute value to ensure that:

- o The host and port number match the current supplier.

- o The database generationID matches the current supplier.
- o The last_change number is smaller than the last change recorded in the supplier's change log. The change log is a log of all the changes made to the supplier's entries. Among other things, this log contains version identification used for synchronization purposes. For more information about configuring the change log for SIR, see "[Configuring the Change Log for SIR](#)".

If the host and port do not match the current supplier, the supplier aborts synchronization and returns an error. This prevents any one replicated entry on the consumer from having multiple suppliers.

6. If no changes are required, the supplier terminates synchronization normally. If changes are required, the supplier updates and/or deletes all appropriate entries on the consumer as indicated by the change log. The supplier also records the last update number applied to the consumer and sets the copiedFrom attribute at the top of the replicated tree on the consumer server. This identifies the tree as being a replica and, more importantly, identifies the supplier server as the master of the information in that tree. The supplier then exits synchronization normally.

CIR Algorithm

If you are using consumer-initiated replication, it is the responsibility of the consumer server to request updates from the supplier server. This process occurs as follows:

1. Based on a schedule you set (which is stored in the consumer server's machine data tree), the consumer server determines that it wants to be updated and binds to the supplier server.
2. The consumer server obtains from its own replication agreements the location of each of its own directory trees supplied to it from another server. The consumer then examines the root point of each of these trees to make sure that a copiedFrom attribute is set on the root point.

If no copiedFrom attribute exists on the tree, the consumer server adds one with the correct information so that all write operations are appropriately referred to the supplier server. This attribute is maintained by the replication subsystem.

The syntax for the copiedFrom attribute is as follows:

```
copiedFrom: host:port generationID last_change
```

For example:

```
copiedFrom: dir.srius.com:389 019980610154028 12
```

where host:port is the host name and port number of the supplier server, generationID is a timestamp generated for the supplier database when the

database is created, and `last_change` is a number generated by the supplier server that increases by one for every change replayed to the consumer. The `generationID` and `last_change` number are only reset if the supplier database is reloaded. Do not modify the `generationID` or the `last_change` number manually. If you do, replication will fail and you will have to reinitialize your consumer servers.

3. The consumer server searches the supplier's change log directory, and compares the contents to its own `copiedFrom` attribute value to ensure:
 - o The host and port number match the current supplier.
 - o The database `generationID` matches the current supplier's generation ID number.
 - o The `last_change` number is smaller than the last change recorded in the supplier's change log. The change log is a log of all the changes made to the supplier's entries. This log contains information used for synchronization purposes. For more information about configuring the change log for CIR, see ["Configuring the Change Log for CIR"](#).

If the host and port do not match the current supplier, the consumer aborts synchronization and returns an error. This prevents any one replicated entry on the consumer from having multiple suppliers.

4. The consumer then checks to see if the last update the consumer recorded in its directory is still contained in the supplier's change log. If it is not, then the consumer has no way of knowing what other changes may have occurred on the supplier since the last time the consumer was updated. In this situation, if `orcauto` is enabled, the consumer reinitializes itself from the supplier server; otherwise, you need to reinitialize the consumer. See ["Initializing Consumers"](#) for more information.
5. If the consumer's last update is still contained in the supplier's change log, then the consumer determines if any changes occurred on the supplier server that must be made to the consumer server's directory. If no changes are required, the consumer terminates synchronization normally. Otherwise, the consumer updates and/or deletes all appropriate entries in its directory tree as indicated by the change log. The consumer sets the appropriate version identification at the same time and then exits synchronization normally.

Machine data

By default, your database actually contains multiple directory trees. One of these trees is used to contain machine data. The machine data tree contains two top-level entries that identify the local server. The first entry uses the `NetscapeMachineData` object class to identify the domain components of the machine on which the server is installed. The second entry uses the `LDAPServer` object class to identify the port on which the LDAP server is listening.

The machine data tree also contains zero or more entries that identify consumer or

supplier servers. On the supplier server of an SIR agreement, the machine data tree contains an entry for each consumer server to which the server replicates data. These consumer server entries use object class `LDAPReplica`. On the consumer server of a CIR agreement, the machine data tree contains an entry for each supplier server. The supplier server entries use object class `cirReplicaSource`.

See the Schema Reference Guide for information on the `NetscapeMachineData`, `LDAPServer`, `LDAPReplica`, and `cirReplicaSource` object classes.

The suffix for the machine data directory tree is

`dc=<serverID>, dc=<domain>, dc=<domain_type>`

For example, if your directory server is running on `directory.airius.com`, then the machine data suffix is

`dc=directory, dc=airius, dc=com`



© Copyright 1999 Netscape Communications Corporation, a subsidiary of America Online, Inc. All Rights Reserved.